arXiv:1405.0762v1 [cs.CG] 5 May 2014

Finding a Curve in a Point Set

Paul Accisano and Alper Üngör

Dept. of Computer & Info. Sci. & Eng.
University of Florida, Gainesville, FL, USA
{accisano, ungor}@cise.ufl.edu

Abstract. Let P be a polygonal curve in $\mathbb{R}^D$ of length n, and S be a point set of size k. The Curve/Point Set Matching problem consists of finding a polygonal curve Q on S such that its Fréchet distance from P is less than a given ε. In this paper, we consider this problem with the added freedom to transform the input curve P by translating it, rotating it, or applying an arbitrary affine transform. We present exact and approximation algorithms for several variations of this problem.

1 Introduction

Matching a curve and a set of points is a typical geometric problem that arises in many science and engineering fields, such as computer aided design, computer graphics and vision, and protein structure prediction [8,11]. In these applications, data is typically gathered as a point set through a scanner, and the goal is often to find certain objects, described as polygonal curves, in the scene. In order to perform a matching, one needs a similarity metric between geometric constructs. In this paper, we study the problem of curve and point set matching, using the Fréchet distance as the similarity metric.

Given a point set and a polygonal curve, the goal is to connect the points into a new polygonal curve that is similar to the given curve. An important factor is whether or not the input curve is allowed to be translated or transformed. The problem in which the curve is fixed in place has been well-studied in the literature [1, 12, 14], and we refer to it in this paper as the ***Curve/Point Set Matching (CPSM)*** problem. Formally, given a polygonal curve P of length n, a point set S of size k, and a real number $\varepsilon > 0$, determine whether there exists a polygonal curve Q on a subset of the points of S such that $\delta_{\mathcal{F}}(P, Q) \leq \varepsilon$.

		Discrete	Continuous
Subset	Unique	NP-C [14]	NP-C [2]
	Non-Unique	P [14]	P [12]
All-Points	Unique	NP-C [14]	NP-C [1]
	Non-Unique	P [14]	NP-C [1]

Table 1: Eight versions of the CPSM problem and their complexity classes.

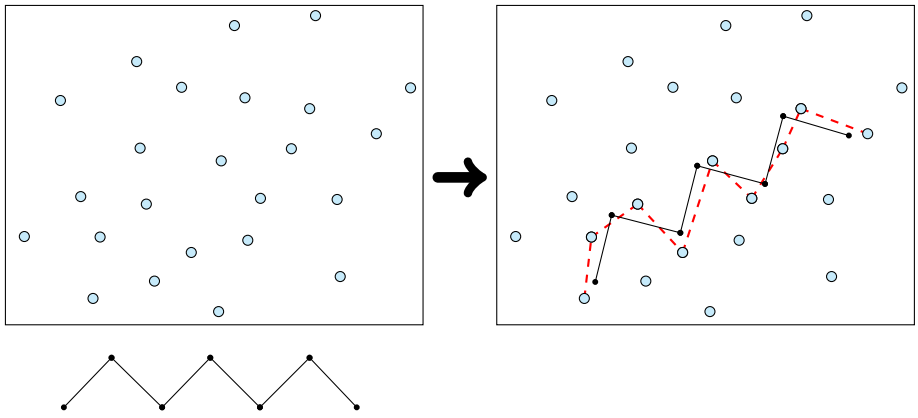

Fig. 1: Example TCPSM instance and solution.

Eight versions of the original CPSM problem can be classified based on whether the use of all points is enforced, whether points are allowed to be visited more than once, and whether the Fréchet distance metric used is discrete or continuous. Table 1 summarizes the versions and their known complexity classes.

Fréchet distance is a powerful and useful metric, but it is very sensitive to positional and rotational differences. If the goal is to locate a curve in a point set that is similar to a given curve, many applications would want the shape of the curve to be the only relevant factor, not its position or orientation. One way to achieve this is to allow the curve to be transformed by an affine transformation from a user-specified set. Indeed, in the literature, the problem of matching two curves under a specified set of affine transformations has been well-studied [11,13]. In this paper, we introduce the ***Transformed Curve/Point Set Matching (TCPSM)*** problem, in which the goal is to find the transformation of P (limited to a user-specified set) that puts it closest to some curve whose vertices are in S. Figure 1 shows an example instance. We note that, in this paper, we focus on the Non-unique versions only, and we use the convention that all the TCPSM problems discussed henceforth refer to the Non-unique version.

Our results. We present an algorithm that makes use of the results in [12] and [13] to solve the Continuous Non-Unique TCPSM. The algorithm runs in $O((nk)^{2d+1}k)$ time for the decision version of the problem, where d is the number of degrees of freedom in the affine transform matrix. Similarly, we show how the Discrete versions are solvable in a similar fashion, resulting in $O((nk)^{d+1})$ time algorithms for both the Subset and All-Points versions. Finally, we give a new algorithm designed from scratch to solve the special case of translations in $\mathbb{R}^2$ for the Discrete versions. Our algorithm runs in $O(n^2k^2\log(nk))$ time, which is an improvement over the general case. Interestingly, this is faster than the best known algorithm for finding the optimal translation when both curves are given. The optimization version of all the problems discussed can be solved with an additional $O(\log(nk))$ factor via parametric search.

2 Preliminaries

Below, we present the notation that will be used throughout the paper, some of which is similar to the notation used by earlier work [1, 4, 12]. More will be introduced later as needed. Given two curves $P, Q : [0, 1] \to \mathbb{R}^d$, the *Fréchet distance* between P and Q is defined as

$$\delta_{\mathcal{F}}(P, Q) = \inf_{\sigma,\tau} \max_{t\in[0,1]} \|P(\sigma(t)), Q(\tau(t))\|$$

where $\sigma, \tau : [0, 1] \to [0, 1]$ range over all continuous non-decreasing surjective functions [10]. Deciding whether two curves have Fréchet distance less than a given ε can be done in $O(nm)$ time, and finding the actual Fréchet distance can be determined in $O(nm \log(nm))$ time by applying parametric search [5]. The Continuous Subset CPSM, which uses the continuous Fréchet distance, is solvable in $O(nk^2)$ time [12].

Discrete Fréchet distance is a variation of the standard Fréchet distance that only takes into account distance at the curve vertices [9]. For two curves P and Q of lengths n and m respectively, a *paired walk* or *coupling sequence* is a pair of integer sequences $(a_1, b_1), \ldots, (a_k, b_k)$, $k \geq \max(n, m)$, with the properties that $(a_1, b_1) = (1, 1), (a_k, a_k) = (n, m)$, and for all i, $(a_{i+1}, b_{i+1}) \in \{(a_i + 1, b_i), (a_i, b_i + 1), (a_i + 1, b_i + 1)\}$. Let W be the set of all paired walks for P and Q. Then the *discrete Fréchet distance* can be defined as:

$$\delta_{\mathrm{F}} = \min_{(a,b)\in W} \max_i \|P_{a_i}, Q_{b_i}\|$$

Note that $\delta_{\mathcal{F}}(P, Q) \leq \delta_{\mathrm{F}}(P, Q)$. The discrete Fréchet distance can be computed directly, without the need for parametric search, in $O(nm)$ time via a dynamic programming algorithm [9]. Recently, Agarwal et al. [3] presented an algorithm that finds the discrete Fréchet distance in $O(nm \log\log(nm)/\log(nm))$ time, breaking the quadratic barrier for this problem. The best known algorithm for the discrete CPSM is still $O(nk)$, for both decision and optimization versions [14].

A polygonal curve P is defined by a set of vertices $P_0, \ldots, P_n$. We use $\overline{P_i}$ to refer to the segment of P between P_{i-1} and P_i, and we use $V(P)$ to denote the set of vertices of P. Let P and Q be curves and $\varepsilon = \delta_{\mathcal{F}}(P, Q)$. A pair of points $(p, q) \in (P \times Q)$, residing in segments $\overline{P_i}$ and $\overline{Q_j}$ respectively, is said to be *feasible* if the subcurves $(P_0, \ldots, P_{i-1}, p)$ and $(Q_0, \ldots, Q_{j-1}, q)$ have Fréchet distance at most ε. If $(p, q) \in (\overline{P_i} \times Q)$ is feasible, we say that q is *visited at* segment i of P. For every point $q \in Q$, there must be at least one feasible pair (p, q), and thus every point in Q is visited in at least one segment of P.

3 Previous Work

3.1 Maheshwari's Algorithm

In [12], Maheshwari et al. gave a dynamic programming algorithm to solve the Continuous Subset CPSM in polynomial time. The algorithm relies on a previous

result by Alt et al. [4] to compute the reachability information for every pair of points in S with respect to P. Starting from the first vertex of P, the algorithm works by computing which points are reachable at which segments, propagating this reachability information through the curve. Once the final vertex of P is reached, if any points in S are still reachable, the algorithm returns true. It runs in $O(nk^2)$ time. We need not modify Maheshwari's algorithm for our purposes; we simply call it as a subroutine.

3.2 Wenk's Algorithm

Wenk [13] looked at a problem similar to the TCPSM. In that problem, instead of matching a curve to a point set, they match a curve to another curve, allowing one to be transformed in order to find the optimal Fréchet distance. As we will make use of this algorithm, we give a full overview in this section.

Let $\mathcal{T}$ be a rationally parameterized[1] subset of the set of affine transformations operating in $\mathbb{R}^D$. Each transformation in $\mathcal{T}$ can be represented by a $D \times D$ linear transformation matrix A and a translation vector $t \in \mathbb{R}^D$. Thus, the number of degrees of freedom d is at most $D^2 + D$, and the parameter space of $\mathcal{T}$ can be represented by $\mathbb{R}^d$. In [13], Wenk considers the problem of finding the $x \in \mathbb{R}^d$ that minimizes $\delta_{\mathcal{F}}(A_x P + t_x, Q)$, where P and Q are polygonal curves. In keeping with many other Fréchet results, the strategy employed is to first solve the decision version of the problem, which asks if there exists an $x \in \mathbb{R}^d$ for which $\delta_{\mathcal{F}}(A_x P + t_x, Q) \leq \varepsilon$, and then use parametric search to find the minimum. For brevity, we let $\tau_x(P) = A_x P + t_x$.

Lemma 1. [13] *For any two polygonal curves P and Q, $\delta_{\mathcal{F}}(\tau_x(P), Q)$ is continuous as a function of $x \in \mathbb{R}^d$.*

Corollary 1. [13] *For a given $\varepsilon > 0$, if there exists some $x \in \mathbb{R}^d$ for which $\delta_{\mathcal{F}}(\tau_x(P), Q) < \varepsilon$, then there exists some $x \in \mathbb{R}^d$ for which $\delta_{\mathcal{F}}(\tau_x(P), Q) = \varepsilon$*

Lemma 2. [13] *If $\delta_{\mathcal{F}}(P, Q) = \varepsilon$, then either:*

- *there exists a vertex x from one curve and a segment $\overline{Z}$ from the other curve for which $\min_{z \in \overline{Z}} \|x, z\| = \varepsilon$,*
- *or there exist two vertices x and y from one curve and a segment $\overline{Z}$ from the other curve for which $\|x, z\| = \|y, z\| = \varepsilon$ for some $z \in \overline{Z}$.*

The above lemma shows that, in order to find a transformation of P which puts its Fréchet distance from Q at exactly ε, we need only consider those transformations which cause one of these two conditions to arise. Both involve a segment and at most two vertices. This leads to the introduction of configurations.

Definition 1. [13] *A* configuration *is a triple $(x, y, \overline{Z})$, where x and y are vertices from one curve (possibly the same vertex) and $\overline{Z}$ is a segment from the other.*

[1] See [13] for a full definition of "rationally parameterized." The set of rationally parameterized affine transforms includes the set of all commonly used transformation types, such as translations, rotations, similarities, and arbitrary affine transforms.

Definition 2. [13] *For a given configuration $c = (P_i, P_j, \overline{Q_k})$, the set of* critical transformations $T^{\varepsilon}_{\text{crit}}(c)$ *is defined as:*

$$T^{\varepsilon}_{\text{crit}}(c) = \begin{cases} \{x \in \mathbb{R}^d \mid \min_{z \in \overline{Q_k}} \|\tau_x(P_i), z\| = \varepsilon\} & : i = j \\ \{x \in \mathbb{R}^d \mid \exists z \in \overline{Q_k} \ \|\tau_x(P_i), z\| = \|\tau_x(P_j), z\| = \varepsilon\} & : i \neq j \end{cases}$$

Symmetrically, if $c = (Q_i, Q_j, \overline{P_k})$,

$$T^{\varepsilon}_{\text{crit}}(c) = \begin{cases} \{x \in \mathbb{R}^d \mid \min_{z \in \tau_x(\overline{P_k})} \|Q_i, z\| = \varepsilon\} & : i = j \\ \{x \in \mathbb{R}^d \mid \exists z \in \tau_x(\overline{P_k}) \ \|Q_i, z\| = \|Q_j, z\| = \varepsilon\} & : i \neq j \end{cases}$$

A transformation τ_x is said to be critical *if it is critical for some configuration c. The arrangement in $\mathbb{R}^d$ of all critical transformations is denoted by $A^{\varepsilon}_{\text{crit}}$.*

Lemma 3. [13] *If there exists a transformation τ_x such that $\delta_{\mathcal{F}}(\tau_x(P), Q) < \varepsilon$ then there exists some face $F \in A^{\varepsilon}_{\text{crit}}$ such that $\delta_{\mathcal{F}}(\tau_y(P), Q) \leq \varepsilon$ for all $y \in F$.*

After showing these lemmas, the strategy employed by Wenk [13] is to obtain a sample point from each face in $A^{\varepsilon}_{\text{crit}}$ and check the Fréchet distance for each corresponding transformation. To obtain such a sample, a result of Basu, Pollack, and Roy [7] is employed. This result shows how to obtain a sample point from each face of an arrangement of semi-algebraic sets in d-dimensional space in $O(M^d)$ time and space, where M is the number of sets. The resulting sample set is termed a *semi-algebraic sample*. Of course, this does require one last lemma.

Lemma 4. [13] *For any configuration c, $T^{\varepsilon}_{\text{crit}}(c)$ is semi-algebraic.*

Since there are a total of n^2m+nm^2 possible configurations, and since deciding if the Fréchet distance of two curves is less than ε takes $O(nm)$ time [5], the total complexity of Wenk's algorithm is $O((nm)^{d+1}(n+m)^d)$.

4 Exact Algorithms for the TCPSM

4.1 Continuous Subset versions

Wenk's algorithm relies on the concept of a configuration: a triple of two vertices from one curve and a segment from the other. The important observation, however, is that each configuration is independent of the others, and the order of the vertices is not relevant to the construction of the arrangement. Therefore, to solve the TCPSM in which one of the curves is unknown, computing a superset of the configurations of a valid curve is sufficient. We can do this by considering all points and potential edges in S. The configurations corresponding to a valid curve Q, if there is one, will be among them. The addition of extra configurations to the arrangement does not affect the correctness of the algorithm.

Since there are k^2 potential edges in the point set, the number of configurations M is $O(nk^2 + n^2k^2) = O(n^2k^2)$. The semi-algebraic sample can be computed in $O(M^d) = O((nk)^{2d})$. Finally, each possible transformation must be checked with Maheshwari's algorithm, which takes $O(nk^2)$ time [12]. Thus, the total running time for the decision version is $O((nk)^{2d+1}k)$. The optimization version can be solved using parametric search, which adds an additional log factor, leading to a running time of $O((nk)^{2d+1}k\log(nk))$.

4.2 Discrete Subset and All-points versions

The Discrete versions of the TCPSM can be solved in a similar fashion, but the lemmas above need to be reexamined to ensure they hold for the discrete Fréchet distance. Lemma 1 trivially holds from the fact that the discrete Fréchet distance is defined by a set of minimums and maximums of Euclidean distance functions.

Lemma 5. *For any two polygonal curves P and Q, $\delta_{\mathrm{F}}(\tau_x(P), Q)$ is continuous as a function of $x \in \mathbb{R}^d$.*

Lemma 2 is not so easy; it only holds true for the continuous Fréchet distance. Its discrete counterpart is as follows.

Lemma 6. *If $\delta_{\mathrm{F}}(P, Q) = \varepsilon$, then there exist two vertices P_i and Q_j for which $\|P_i, Q_j\| = \varepsilon$.*

Proof. The proof is much simpler that the corresponding continuous Lemma, owing to the fact that the number of unique paired walks is finite. Since the expression inside the minimum and maximum of the discrete Fréchet distances' definition is the distance between two vertices, δ_{F} must be one of these distances. □

Thus, configurations for the discrete problem are redefined as a pair (P_i, Q_j), and critical transformations of a configuration $c = (P_i, Q_j)$ are redefined as:

$$T^{\varepsilon}_{\mathrm{crit}}(c) = \{x \in \mathbb{R}^d \mid \|\tau_x(P_i), Q_j\| = \varepsilon\}$$

The discrete Fréchet equivalent of Lemma 4 trivially holds for this new definition, as $T^{\varepsilon}_{\mathrm{crit}}(c)$ is defined by a simple polynomial expression. The discrete Fréchet equivalent of Lemma 3 also holds, but although the proof is very similar to proof for the continuous version given in [13], we feel it is different enough to warrant at least a sketch of the proof, highlighting the differences between the two.

Lemma 7. *If there exists a transformation τ_x such that $\delta_{\mathrm{F}}(\tau_x(P), Q) < \varepsilon$ then there exists some face $F \in A^{\varepsilon}_{\mathrm{crit}}$ such that $\delta_{\mathrm{F}}(\tau_y(P), Q) \leq \varepsilon$ for all $y \in F$.*

Proof. By the continuity of discrete Fréchet distance and its corollary, the existence of a transformation τ_x for which $\delta_{\mathrm{F}}(\tau_x(P), Q) < \varepsilon$ implies the existence of a transformation $\tau_{x^=}$ for which $\delta_{\mathrm{F}}(\tau_{x^=}(P), Q) = \varepsilon$. By Lemma 6, there is some configuration c for which $x^= \in T^{\varepsilon}_{\mathrm{crit}}(c)$. Let F be the connected component of $T^{\varepsilon}_{\mathrm{crit}}(c)$ that contains $x^=$. If $\delta_{\mathrm{F}}(\tau_y(P), Q) \leq \varepsilon$ for all $y \in F$, the claim is shown. Otherwise, let $x^> \in F$ be such that $\delta_{\mathrm{F}}(\tau_{x^>}(P), Q) > \varepsilon$.

Let R be a curve on F such that $R(0) = x^=$ and $R(1) = x^>$, and let $\varepsilon(r) = \delta_{\mathrm{F}}(\tau_r(P), Q)$ for $r \in R$. Assume without loss of generality that $\varepsilon(R(s)) > \varepsilon$ for any $s > 0$. From the definition of discrete Fréchet distance and the continuity of critical transformations as a function of ε, we have that there must exist some open neighborhood S around 0 for which $R(s) \in T^{\varepsilon(R(s))}_{\mathrm{crit}}(c^*)$ for all $s \in S$ and some configuration c^*. Therefore, since $T^{a}_{\mathrm{crit}}(c^*) \cap T^{b}_{\mathrm{crit}}(c^*) = \emptyset$ for $a \neq b$, we have

that $R(0) \in T^{\varepsilon}_{\text{crit}}(c^*)$, but $R(s) \notin T^{\varepsilon}_{\text{crit}}(c^*)$ for any small value of s. However, $R \subseteq F \subseteq T^{\varepsilon}_{\text{crit}}(c)$. Therefore, $c \neq c^*$ and $x^{=} \in T^{\varepsilon}_{\text{crit}}(c) \cap T^{\varepsilon}_{\text{crit}}(c^*)$. We then apply the same argument to the lower dimensional face $F \subseteq T^{\varepsilon}_{\text{crit}}(c) \cap T^{\varepsilon}_{\text{crit}}(c^*)$ of $A^{\varepsilon}_{\text{crit}}$, and the claim follows by induction.

□

With the equivalent lemmas in hand, the solution approach remains the same. However, the time complexity is of course much faster than the continuous version, owing to the smaller number of configurations and the simpler algorithm for the discrete CPSM. The number of configurations is $O(nk)$, and each sample point in the transformation space can be tested in $O(nk)$ time, leading to a final running time of $O((nk)^{d+1})$ for the decision version and $O((nk)^{d+1}\log(nk))$ for the optimization version. The parametric search analysis in [13] for computing the ε values to check applies straightforwardly.

5 Discrete CPSM for Translations in $\mathbb{R}^2$

The exact algorithms presented in the previous section all rely on a combination of previous algorithms in the literature. However, for the special case of translations in $\mathbb{R}^2$ for the Discrete TCPSM, we can use a different approach and solve the problem directly, without resorting to either Wenk's algorithm or Maheshwari's. We refer to the version of the TCPSM where the set of allowed transformations is restricted to only translations as the **tCPSM**. While the algorithm presented in Section 4 can solve this version of the problem in $O(n^3k^3)$ time, the algorithm presented below can solve it in $(n^2k^2\log(nk))$ time.

The Discrete Subset tCPSM can be reformulated as follows: Find a translation of P such that every vertex of P is within ε of some point in S. Consider the set of all translations $t \in \mathbb{R}^2$ that put a vertex P_i within ε of a point $s \in S$. This set is a disk in the plane. With all combinations of vertices of P and points in S, we have $O(nk)$ such disks. If we consider each disk to be "colored" with a unique color corresponding to its associated vertex of P, the problem can be reformulated as follows: Find a single point that lies within at least one disk of every color. This problem can be solved by a plane-sweep algorithm, which we now describe.

The event points of the plane-sweep algorithm are the tops and bottoms of each of the $O(nk)$ disks, as well as their $O(n^2k^2)$ intersection points. The intersection of the sweep line with each disk divides the sweep line into intervals. For each interval, we store a membership array that keeps track of how many disks of each color that interval is inside, as well as a counter that records how many elements of the membership array are non-zero. When the sweep line reaches the top of a new disk, the interval in which the disk top resides is split into three pieces, each with the copy of the membership array. The array element of the color of the new disk is incremented in the middle piece's array, which corresponds to the sweep line interval inside the new disk. If that element was zero, the non-zero counter is incremented. When the sweep line reaches the bottom

of a disk, the corresponding interval is removed and the two adjacent intervals are merged; their membership arrays will be identical. At an intersection point, the membership arrays and non-zero counters of the three intervals involved are updated accordingly. If at any time some interval's non-zero counter becomes equal to n, the algorithm stops and returns a point of that interval as the solution. Otherwise, if the sweep line passes the last disk, the algorithm reports that there is no solution. The full pseudo-code listing is shown Algorithm 1.

Algorithm 1 Discrete tCPSM in $\mathbb{R}^2$ (P, S, ε)

```
1: Compute the disks corresponding to P and S and their intersection points.
2: Create a priority queue using the intersection points and the top and bottom
   of each disk.
3: Using a balanced binary search tree, initialize a sweep line with a single
   interval I, whose membership array M_I and non-zero counter Z_I are zeroed.
4: while the queue is non-empty do
5:     Dequeue an event point p.
6:     Find the sweep line interval I in which p resides.
7:     if p is the top of a disk with color c then
8:         Split I into I, J, K, copying its membership array.
9:         Increment M_J[c].
10:        Increment Z_J if M_J[c] was zero.
11:        if Z_J = n then
12:            return p as the solution.
13:        end if
14:    else if p is the bottom of a disk then
15:        Delete I from the sweepline.
16:        Merge its two neighbors, deleting the membership array of one.
17:    else if p is the intersection point of two disks of colors c_1 and c_2 then
18:        Increment or decrement M_I[c_1] and M_I[c_2] accordingly.
19:        Increment or decrement Z_I accordingly.
20:        if Z_I = n then
21:            return p as the solution.
22:        end if
23:    end if
24: end while
25: return that there is no solution.
```

5.1 Analysis

Lines 1 through 3 take $O(n^2k^2)$ time. The while loop of line 4 iterates over $O(n^2k^2)$ points, taking $O(\log(nk))$ time for each dequeue operation. Using a balanced binary search tree as the sweep line data structure allows lines 6 and

15 to be accomplished in $O(\log(nk))$ time. Line 8 takes $O(n)$ time, but is only executed for disk top points, of which there are only $O(nk)$. The rest of the lines in the while loop are all $O(1)$. Thus, the total running time of the algorithm is $O(n^2k^2\log(nk))$. Surprisingly, this is faster than the best known algorithm to solve the problem when both curves are known, which is $O(n^3m^3)$ [11].

5.2 All-points Version

If we add the additional constraint that every point in S be within ε from some vertex in p, the problem becomes equivalent to the All-Points variation. To solve this variation, we define each disk to be of two colors, one for its corresponding vertex in P and one for its corresponding point in S. The algorithm is easily modified to increment or decrement two entries in the membership array instead of one when a new disk or a disk intersection point is encountered. The non-zero counter of an interval will have to rise to $n + k$ to report a valid solution, and copying the membership array in line 8 will take $O(n + k)$ time. However, the time complexity is still identical, as the cost of iterating through all event points in sorted order still dominates.

6 Approximation Algorithms

6.1 3-Approximation for Continuous All-Points TCPSM

In the appendix of [2], we presented a 3-factor approximation algorithm for the Continuous All-points version of the CPSM, which is NP-complete. The approximation algorithm works by first deciding if there is curve that, in addition to having Fréchet distance at most ε from P, visits each point in S at its closest segment. We call such a curve *NS-compliant*.

Definition 3. *Let $\varepsilon = \delta_{\mathcal{F}}(P, Q)$, where P and Q are curves, and for a given point q, let $c(q)$ be the index of the segment of P nearest q. Q is said to be* NS-compliant *if, for every point q in the vertex set $V(Q)$, there exists a vertex $Q_i = q$ that is visited at $\overline{P_{c(q)}}$.*

Lemma 8. *Let $\varepsilon = \delta_{\mathcal{F}}(P, Q)$, where P and Q are curves. There exists an NS-compliant curve Q' with the same vertex set as Q such that $\delta_{\mathcal{F}}(P, Q') \leq 3\varepsilon$.*

Proof. Since Q might not be NS-compliant, there may be some vertices of Q that are not visited at their closest segment of P; let s be such a point. Since s is visited at a segment other than its closest, it can be no further than ε away from $\overline{P_c(s)}$. Let p be a point in $\overline{P_c(s)}$ that is within ε of s. Recall that there is always at least one feasible pair for any point on P; let (s', p) be feasible. Then, add s' as a new vertex of Q. Note that the distance between s and s' is at most 2ε. Repeating this process for every point not visited at its closest segment yields a new curve Q^*. Since each new vertex has been added along an existing segment, $\delta_{\mathcal{F}}(P, Q) = \delta_{\mathcal{F}}(P, Q^*)$.

Now, merge each s' with its corresponding s by translating the former to the position of the latter, yielding a new curve Q' with a potentially different Fréchet distance from P. Let σ and τ be reparameterizations of P and Q^*, and consider the point $Q^*(\tau(t))$ for some $t \in [0,1]$, which lies on some segment of Q^*. The endpoints of the corresponding segment in Q' may have been displaced up to 2ε, and thus the point $Q'(\tau(t))$ may be up to 2ε away from $Q^*(\tau(t))$. Therefore, $\|P(\sigma(t)), Q'(\tau(t))\|$ can be at most 2ε larger than $\|P(\sigma(t)), Q^*(\tau(t))\|$. Finally, since the Fréchet distance is the infimum of the maximum distance over all reparameterizations, we have that $\delta_{\mathcal{F}}(P, Q') \leq \delta_{\mathcal{F}}(P, Q^*) + 2\varepsilon = 3\varepsilon$. □

An algorithm presented in the appendix of [2] can be used for deciding whether there exists a curve which visits all points in S, has Fréchet distance at most ε from P, and is NS-compliant. It runs in $O(nk^2)$ time, and the optimal NS-compliant curve can be found in $O(nk^2 \log(nk))$ time by way of parametric search. Since the optimal curve can be made NS-compliant while only increasing its Fréchet distance by a factor of 3, this yields a 3-approximation algorithm.

Theorem 1. *The Continuous All-points CPSM can be 3-approximated in $O(nk^2)$ time.*

When integrating this result into Wenk's framework, it is tempting to simply apply the decision algorithm to check for an NS-compliant curve at each semi-algebraic sample point in the transformation space, and then use parametric search to find the optimal NS-compliant curve. However, this will not work. As P is translated or transformed, the closest segment to each point in S can change, which means curves that were NS-compliant for one transformation may be non-compliant for others. Therefore, even if $\tau(P)$ is exactly the optimal Fréchet distance away from the optimal NS-compliant curve, the curve may not be NS-compliant for that particular value of τ, causing the decision algorithm to return FALSE.

The solution to this problem is to remember that our goal is not to find the optimal NS-compliant curve, but to find the optimal unrestricted curve. To that end, we modify the algorithm so that, at each step ε of the parametric search, it checks each sample point in the transformation space for an NS-compliant curve with Fréchet distance at most 3ε.

Let Q_{opt} be the optimal unrestricted curve that visits every point in S, and let ε_{opt} be the Fréchet distance from Q_{opt} to the optimal transformation of P. As before, the configurations corresponding to Q_{opt} will be among those added to the arrangement, and the others will have no effect on the correctness. For a given parametric search step $\varepsilon \geq \varepsilon_{\text{opt}}$, there will be at least one sample point τ in the parameter space for which $\delta_{\mathcal{F}}(\tau(P), Q_{\text{opt}}) = \varepsilon$. For this value of τ, Lemma 8 guarantees that there will be an NS-compliant curve with Fréchet distance at most 3ε from $\tau(P)$, regardless of which point happens to be closest to which segment. Thus, the parametric search will continue downward and is guaranteed to terminate at some step ε for which $\varepsilon_{\text{opt}}/3 \leq \varepsilon \leq \varepsilon_{\text{opt}}$, yielding a transformation τ and a curve Q with $\varepsilon_{\text{opt}} \leq \delta_{\mathcal{F}}(\tau(P), Q) \leq 3\varepsilon_{\text{opt}}$.

Theorem 2. *The Continuous All-Points TCPSM can be 3-approximated in* $O((nk)^{2d+1}k\log(nk))$ *time.*

6.2 (1+ε)-Approximation

The running times of the exact algorithms discussed in the previous sections are all quite high, even when the degrees of freedom are few. Because of this, it makes sense to look at approximation algorithms that might have lower time complexity. In [6], the authors present a $(1+\varepsilon)$ approximation algorithm for the two-curve matching problem under translations in $\mathbb{R}^2$. In this section, we generalize their approach to work for the tCPSM problem in $\mathbb{R}^d$. We make use of the following Lemma, which was proven in [13], and a key observation about both types of Fréchet distance.

Lemma 9. [13] *Let P and Q be curves and let t be a translation vector in $\mathbb{R}^d$. Then $\delta_{\mathcal{F}}(P+t,Q) \leq \delta_{\mathcal{F}}(P,Q) + \|t\|$. The same applies to discrete Fréchet distance.*

Observation 3. *Let P and Q be curves. Then $\|P_0, Q_0\| \leq \delta_{\mathcal{F}}(P,Q) \leq \delta_{\mathrm{F}}(P,Q)$.*

This suggests the strategy of translating P such that its start point overlaps with the start point of the optimal curve Q. In fact, this strategy guarantees a Fréchet distance from Q at most twice the optimal. To show this, let t_{opt} be the optimal translation of P and let $t_{\mathrm{apx}} = Q_0 - P_0$ be the translation that, when applied to P, lines up the two initial points of the curve. By the observation above, we have that $\|t_{\mathrm{apx}} - t_{\mathrm{opt}}\| = \|P_0 + t_{\mathrm{opt}}, P_0 + t_{\mathrm{apx}}\| = \|P_0 + t_{\mathrm{opt}}, Q_0\| \leq \delta_{\mathcal{F}}(P + t_{\mathrm{opt}}, Q)$. By Lemma 9, this shows that $\delta_{\mathcal{F}}(P + t_{\mathrm{apx}}, Q) = \delta_{\mathcal{F}}(P + t_{\mathrm{opt}} + (t_{\mathrm{apx}} - t_{\mathrm{opt}}), Q) \leq 2\delta_{\mathcal{F}}(P + t_{\mathrm{opt}}, Q)$. Finally, we can improve this 2-approximation to a $1+\varepsilon$ approximation by centering an ε-lattice of width 2 around Q_0 and trying every lattice point.

Of course, all this assumes that the optimal curve Q is known, which it is not. However, its start point must be one of the k input points. Trying each point in S adds another factor of k, leading to a final worst case running time of $O(nk^3\log(nk)/\varepsilon^d)$ for the continuous versions and $O(nk^2/\varepsilon^d)$ for the discrete. We note that, in practice, this additional factor of k will typically have a very small associated constant, since those translations that do not also put the end point of P close to a point in S can quickly be ruled out.

Theorem 4. *The tCPSM can be $(1+\varepsilon)$-approximated in $O(nk^3\log(nk)/\varepsilon^d)$ time for the Continuous Subset version and $O(nk^2/\varepsilon^d)$ for the Discrete Subset and Discrete All-points versions.*

In addition, the same idea can be applied to the 3-approximation algorithm for the Continuous All-points CPSM discussed in Section 6.1 to yield a $(3+\varepsilon)$ approximation for the corresponding tCPSM version.

Corollary 2. *The Continuous All-points tCPSM can be $(3+\varepsilon)$-approximated in $O(nk^3\log(nk)/\varepsilon^d)$.*

7 Conclusion

In summary, we have shown how various versions of the TCPSM can be solved in polynomial time. The Continuous Subset version can be solved exactly and the Continuous All-Points version can be 3-approximated, both in $O((nk)^{2d+1}k\log(nk)$ time. Both Discrete versions can be solved in $O((nk)^{d+1}\log(nk)$ time, and for the special case of translations in $\mathbb{R}^2$, $O(n^2k^2\log^2(nk))$ time. Furthermore, a $(1+\varepsilon)$ approximation algorithm exists for translations in $\mathbb{R}^d$ that runs in $O(nk^3\log(nk))$ time.

References

1. P. Accisano and A. Üngör. Hardness results on curve/point set matching with Fréchet distance. In *Proc. of the 29th European Workshop on Comp. Geom.*, 2013.
2. P. Accisano and A. Üngör. Matching curves to imprecise point sets using Fréchet distance. *arXiv preprint arxiv.org*, 2014.
3. Pankaj K Agarwal, Rinat Ben Avraham, Haim Kaplan, and Micha Sharir. Computing the discrete fréchet distance in subquadratic time. *arXiv preprint arXiv:1204.5333*, 2012.
4. H. Alt, A. Efrat, G. Rote, and C. Wenk. Matching planar maps. In *Proc. of the 14th annual ACM-SIAM Symp. on Discrete Algorithms*, pages 589–598, 2003.
5. H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. *Int. J. of Comp. Geom. & Appl.*, 5(01n02):75–91, 1995.
6. H. Alt, C. Knauer, and C. Wenk. Matching polygonal curves with respect to the Fréchet distance. In *STACS 2001*, pages 63–74. Springer, 2001.
7. S. Basu, R. Pollack, and M. Roy. A new algorithm to find a point in every cell defined by a family of polynomials. In *Quantifier elimination and cylindrical algebraic decomposition*, pages 341–350. Springer, 1998.
8. S. Brakatsoulas, D. Pfoser, R. Salas, and C. Wenk. On map-matching vehicle tracking data. In *Proc. Int. Conf. on Very Large Data Bases*, pages 853–864, 2005.
9. T. Eiter and H. Mannila. Computing discrete Fréchet distance. Technical Report CD-TR 94/65, Christian Doppler Laboratory, 1994.
10. G. Ewing. *Calculus of variations with applications*. Dover Pub., New York, 1985.
11. M. Jiang, X. Ying, and B. Zhu. Protein structure–structure alignment with discrete Fréchet distance. *J. of Bioinfo. and Comp. Biology*, 6(01):51–64, 2008.
12. A. Maheshwari, J. Sack, K. Shahbaz, and H. Zarrabi-Zadeh. Staying close to a curve. In *Canadian Conf. on Comp. Geom.*, pages 55–58, 2011.
13. C. Wenk. *Shape matching in higher dimensions*. PhD thesis, Freie Universität Berlin, Universitätsbibliothek, 2002.
14. T. Wylie. Discretely following a curve. In *Combinatorial Optimization and Applications*, pages 13–24. Springer, 2013.